

00/60/80
S-127

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
O. Box 272400
Fort Collins, Colorado 80528-9599

03-10-06

PATENT APPLICATION

ATTORNEY DOCKET NO. 10992150-1

IN THE U.S. PATENT AND TRADEMARK OFFICE
Patent Application Transmittal Letter

ASSISTANT COMMISSIONER FOR PATENTS
Washington, D.C. 20231

09/5/2000
JC553 U.S. PTO
03/09/00

Sir:

Transmitted herewith for filing under 37 CFR 1.53(b) is a(n): Utility Design

original patent application,
 continuation-in-part application

INVENTOR(S): Tony M. Brewer, et al

TITLE: **Protocol For Insuring Exactly Once Semantics Of Transactions Across An Unordered, Unreliable Network**

Enclosed are:

The Declaration and Power of Attorney. signed unsigned or partially signed
 2 sheets of drawings (one set) Associate Power of Attorney
 Form PTO-1449 Information Disclosure Statement and Form PTO-1449
 Priority document(s) (Other) (fee \$ _____)

CLAIMS AS FILED BY OTHER THAN A SMALL ENTITY				
(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) TOTALS
TOTAL CLAIMS	20 — 20	0	X \$18	\$ 0
INDEPENDENT CLAIMS	3 — 3	0	X \$78	\$ 0
ANY MULTIPLE DEPENDENT CLAIMS	0		\$260	\$ 0
BASIC FEE: Design \$310.00 ; Utility \$690.00				\$ 690
TOTAL FILING FEE				\$ 690
OTHER FEES				\$
TOTAL CHARGES TO DEPOSIT ACCOUNT				\$ 690

Charge \$ 690 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16, 1.17, 1.19, 1.20 and 1.21. A duplicate copy of this sheet is enclosed.

"Express Mail" label no. EL483249633USDate of Deposit 3/9/00

Respectfully submitted,

Tony M. Brewer, et al

By Leslie P. Gehman

Leslie P. Gehman

Attorney/Agent for Applicant(s)
Reg. No. P45,624Date: March 9, 2000

Telephone No.: (970) 898-3642

**PROTOCOL FOR INSURING EXACTLY ONCE SEMANTICS OF TRANSACTIONS
ACROSS AN UNORDERED, UNRELIABLE NETWORK**

Tony M. Brewer
5227 Mariners Drive
Plano, Texas 75093
Citizenship: U.S.A.

Gregory S. Palmer
3012 Mason Drive
Plano, Texas 75025
Citizenship: U.S.A.

TECHNICAL FIELD

The invention relates in general to data transmission over a network and in particular to efficiently monitoring such data transmission over an unreliable network.

BACKGROUND

Generally when conducting communication including transmitting messages, packets, or transactions between data processing or storage devices, some form of protocol or handshaking is required to manage and coordinate the communication. Where only two devices are transmitting information back and forth, such communication management may be relatively straightforward. Generally, direct and continuous contact may be established between two devices enabling misdirected transmissions, omitted information, and erroneous data to be readily identified and corrected.

Such communication management generally becomes more difficult and complex where a plurality of sources communicate with a plurality of destinations across a network as is shown in FIGURE 1. In such a network, direct and continuous contact between a source and destination for a particular message will generally not be available. Instead the communication process will generally require that messages be transmitted via one or more intermediate points in the network in between transmission by the source and reception by the ultimate destination. A number of problems may be introduced at the intermediate points in the network including the loss of information, introduction of erroneous information, or a hardware failure causing a complete transmission failure. Such events may leave both the source and destination of the transmission unaware of the problem and unable to resolve it.

One solution employed in the prior art to deal with this problem is the association of sequence numbers with messages or packets transmitted from a source to a destination. With sequence numbers attached to each message, destination devices can check the sequence number associated with a message to ensure that it is the next message in sequence from the source location concerned. If a duplicate message is received at a destination device, the message is generally ignored. If a sequence of messages is received wherein a sequence number has been skipped, a protocol active between the destination and the source will generally cause the skipped message to be resent.

The use of sequence numbers for communication management in a multiple source, multiple destination network will generally impose a substantial processing burden on the system. Generally, each source-destination pair will have its own set of sequence numbers so as to uniquely identify each source and destination within the system. The total number of sequence number sets will generally increase exponentially with increasing network size. For example, where ten processors communicate with ten memory boards, one hundred sequence number sets may be required. Whereas, for one thousand processors and one thousand memory boards, one million sequence number sets will generally be required in order to uniquely identify each source-destination pair. Such a large number of sequence number sets generally increases the burden on each source and destination by increasing the number of bits required to identify the message being transmitted.

Furthermore, such a system is difficult to effectively scale down. For example, where a system is initially designed to accommodate one million sequence number sets, the system may handle the overhead associated with the sequence numbers in an acceptable manner. However, if the same chips were to be used in a small topology involving, for example, one processor and one memory board, a chip set designed to handle one million possible sequence number sets will be employed to handle just one sequence number set for the small topology, thereby imposing an enormous and unproductive overhead cost to operation of the small topology.

Generally, a typical action responsive to a determination by a destination device that a transmission is faulty, or was lost, is to issue a request to the source to retransmit the pertinent information. In a single source, single destination system, such an approach is generally quite feasible. However, in a multiple source, multiple destination system, the very data which identifies the source of the data could be lost during transmission, making it extremely difficult for the destination device to issue a retransmission request. Accordingly, the use of sequence numbers to identify transactions in a multiple source, multiple destination system may make both detection and correction of a transmission error very difficult.

One alternative to the use of a distinct set of sequence numbers for every source-destination pair in a large network is the use of software to dynamically allocate channels at a processor or other communication component in the network. Generally, software operates to allocate a channel within a processor to communicate with another component in source-destination pair for a period of time. When communication is established with a different component, the channel may be re-allocated for communication with the new component and a set of sequence numbers established on-the-fly. This approach generally avoids the requirement of having an immense number of sets of sequence numbers but generally imposes latency on the communication process. Specifically, the time required to initially allocate and re-allocate channels for communication with different components or devices requires a finite period of time. In networks where communication between devices generally experiences very short latencies, such as connections between chips in a computer system, the time required for channel allocation and re-allocation generally imposes an unacceptable delay.

Another alternative to the option of using distinct sequence number sets for all source-destination device pairs in a network is the use of pre-allocated hardware channels on all participating communication devices. FIGURE 2 depicts destination devices 202 and 203 having dedicated channel connections to processor agent device 201. This approach presents the advantage of not incurring the communication latencies generally caused by the previously discussed software-based dynamic channel allocation approach, but imposes a substantial requirement for additional hardware functionality on each participating communication device. This added hardware requirement will generally impose substantial additional system cost.

The use of sequence numbers to identify transactions generally imposes a forced ordering of transactions between a source and a destination even where such ordering is not required by higher protocol levels. This is because the destination generally only identifies missed packets by receiving a series of packets whose sequence numbers are not in proper order. However, in a multiple source, multiple destination network, packets with higher

sequence numbers may arrive at a destination device prior to an earlier sent packet with a lower sequence number due to the higher numbered packet having taken a more expeditious path through the network. This out-of-order arrival does not necessarily indicate that a packet has been lost. Therefore, a buffer will generally be employed to hold a packet with a higher sequence number so that packets with lower sequence numbers are allowed to pass through the buffer and arrive at the destination first. Such forced ordering thereby generally operates to slow down the transmission of transactions and thereby the network as a whole.

Therefore, it is a problem in the art that the use of distinct sequence number sets for all source-destination pairs in a multiple source, multiple destination network (or large complex network) imposes a burden on communication and computation in the network.

It is a further problem in the art that the use of sequence numbers in a multiple source, multiple destination network may prevent accurate identification of the device which transmitted a data packet identified as defective at a destination device.

It is a still further problem in the art that the use of sequence numbers for packet identification in large complex network generally imposes a need for forced ordering of transactions in said network.

It is a still further problem in the art that the use of software based dynamically allocated channels for sequence numbering of data packets generally imposes unacceptable latency on communication within a computer network.

It is a still further problem in the art that the use of dedicated pre-allocated hardware channels for sequence numbering of data packets in a network generally imposes a substantial cost burden on the system.

SUMMARY OF THE INVENTION

These and other objects, features and technical advantages are achieved by a system and method which enables a storage system employing temporary and reusable entries for monitoring transactions between data sources and destinations across an unordered unreliable transport mechanism. The monitoring of a transaction, and the status thereof, may be performed at a source or destination site. Monitoring of transactions at intermediate points within the transport mechanism is preferably unnecessary with the present invention. Where repetition of the communication of transaction packets or messages is necessary, end to end retry is preferably employed. Generally, end to end retry refers to initiation of a retransmission from one end of a communication link rather than attempting such retransmission from an intermediate point within this link.

In a preferred embodiment, a sequence of transactions is transmitted between a source site, and a destination site, to manage the delivery and acknowledgment of packets between the two sites and to provide unambiguous information regarding the status of all transactions initiated by the source. The inventive mechanism preferably provides exactly once semantics between source and destination over an unordered unreliable network.

In a preferred embodiment, a database, which may be a content addressable memory, associated with each destination device operates as a registry for pending transactions within the network. A preferred embodiment approach involves transmitting a sequence of transaction messages between the source and destination to ensure proper execution of the transaction and to maintain information on the status of the transaction. An exemplary sequence of transaction messages preferably includes a reservation request, a reservation response, a release request, and a release response, which transaction messages are discussed in greater detail in the Detailed Description.

In a preferred embodiment, the transport mechanism is a computer network linking together a plurality of processors, switching devices, memory devices, and other data devices in an unordered and unreliable network. Alternatively, the transport mechanism may include any set of connections between data storage or data processing devices including but not

limited to: local area networks and wide area networks. Herein, "unreliable" refers to a condition wherein a source initiating a transaction may not depend on the pertinent network to transmit the message to the desired destination, and may not assume that the destination will execute the desired transaction.

5 Instead, verification of both the transmission of a request by a source and of the successful completion of a requested transaction by a destination is preferred for the inventive mechanism to conclude that a transaction has been successfully completed. The "end to end" reporting and verification system contemplated herein preferably avoids having to perform enquiries at intermediate points in a network linking the source and destination and removes a need to assume that communication along this network has been successful.

10 Herein, a site which initiates a sequence of transaction messages to effect a transaction at another site is a "source site," and the "other" site at which the transaction is executed is a "destination site." The source and destination sites may each include a single component or may include a plurality of components. The terms "source" and "destination" are generally used synonymously with "source site" and "destination site," respectively. In the Detailed Description, the source agent is a device specifically associated with a processor at a source site and the "destination memory control" is a device specifically associated with a memory board at a destination site. However, the term "agent device" is used herein to refer generally to a component or device which serves as an agent to another device at that site, such as a 20 processor or memory board. Therefore, the term agent device preferably applies to both the "source agent" and the "destination memory control." Since, both the processor and memory board are devices which are being served by other devices, they are referred to generically herein as "master devices." Although the term device generally refers to a single hardware component, the various devices at a site need not be discrete components, but may be 25 integrated into a single package. Furthermore, the function of such devices may optionally be performed in software.

Herein, the term "transaction" includes an operation to be performed at a site remote from source site requesting the operation as well as the communication steps necessary to

effect this operation. An operation may be a memory read, memory write, device control operation, or other data transfer or manipulation operation. A sequence of communication steps is preferably undertaken to complete a transaction, and the information packages transmitted in connection with a transaction are referred to herein as "transaction packets," or "transaction messages." These terms may be further qualified according to the content of each message. For brevity, a transaction which includes a read operation is referred to as a "read transaction," and a transaction which includes a write operation is referred to as a "write transaction."

Herein, the term "exactly once semantics" generally refers to a condition wherein each transaction in the system is preferably acted upon only once. Generally, an error condition exists if a transaction is either not performed at all (i.e. zero occurrences), or performed two or more times.

20

In a preferred embodiment, a source may be a computer processor having an associated agent chip for conducting communication regarding the status of transactions, thereby removing the processor from a critical path for such communication. In a preferred embodiment, the destination is a memory board or other data storage device which preferably has an agent device for handling communication pertaining to transactions with other devices. One possible type of agent device for the data storage device is a destination memory control. Preferably, both the source and destination sites include databases for preserving information regarding the status of a transaction while said transaction is active within the network.

25

In a preferred embodiment, agent devices included in source and destination sites operate to ensure delivery of transactions across an unreliable network. The use of agent devices for such a maintenance function preferably improves the efficiency of the source and destination devices by relieving them of the processing burden of monitoring transaction communication. Preferably, databases are included at the source and destination sites for storing information identifying the transactions active in the network and the status of these transactions. Preferably, data locations in these databases are used temporarily for particular transactions and are infinitely reusable. Preferably, when a new transaction is generated, an

entry associated with that transaction is stored in one or more databases and continues to be used to reference that transaction until the transaction is no longer active in the network. Upon completion of a transaction, the entry for that transaction is preferably released for reuse in association with a new transaction. The data entry mechanism of the present invention is thereby preferably able to serve an infinite stream of successive transactions with a finite number of database entries by releasing device entries upon conclusion of a transaction, and reallocating them upon initiation of a new transaction. Where a transaction, for any reason, cannot be completed, a response packet containing relevant error information is preferably provided to any device awaiting a response in connection with the incomplete transaction. If no device is waiting for a response in connection with an incomplete transaction, error information associated with the transaction is preferably logged locally before resources associated with the incomplete transaction are enabled to be used for other purposes.

In a preferred embodiment, transaction messages may be transmitted by the source sites in any order and received by the destination sites in any order with no adverse effect on the ability of the inventive mechanism to record and manage communication activities between sources and destinations of transactions in a transport mechanism, which may be a computer network. This feature preferably provides for transaction ordering flexibility, thereby removing a constraint on the network generally present when employing sequence numbers. This represents an advantage over the prior art where rigid ordering of transactions generally operates as a constraint on network activity, and in certain cases, causes otherwise executable transactions to remain idle, to enable transactions with lower sequence numbers to be processed first.

In a preferred embodiment of the present invention, the delays associated with 25 allocating a channel for sequence numbers in software at a source or destination site is preferably unnecessary because of the operation of the agent devices and databases at each site as discussed above. Furthermore, the use of temporary reusable entries within the databases for managing the communication of transactions over the network preferably

renders the allocation of dedicated hardware channels within a communication device unnecessary.

In a preferred embodiment, if the source does not receive a reply packet responsive to a packet transmitted by the source in a pre-determined time-out period, the source preferably retransmits its packet. This retransmission presents the possibility of having two identical source-issued messages propagating through the system at the same time. In order to prevent this from occurring, the source includes a time stamp in the outgoing packet and establishes a time-out threshold for the outgoing packet, which is preferably shorter than the time-out threshold for the transmitting source, thereby enabling any device in the network to drop the packet from the network when the packet time-out threshold is reached. The system thereby preferably ensures that there will never be two packets being simultaneously transmitted through the network.

Therefore, it is an advantage of a preferred embodiment of the present invention that the computationally burdensome operation of designating distinct sequence number sets for source-destination pair in a communication network is unnecessary.

It is a further advantage of a preferred embodiment of the present invention that forced ordering of transactions is not required, thereby enabling transactions to be efficiently processed at a destination device in the order in which they are received.

It is a still further advantage of a preferred embodiment of the present invention that pre-allocated hardware channels are not required in order to manage the communication of transactions.

It is a still further advantage of a preferred embodiment of the present invention that greater certainty regarding the status of a transaction in the network is provided than was available using the sequence number identification system of the prior art.

The foregoing has outlined rather broadly the features and technical advantages of the present invention in order that the detailed description of the invention that follows may be better understood. Additional features and advantages of the invention will be described

hereinafter which form the subject of the claims of the invention. It should be appreciated by those skilled in the art that the conception and specific embodiment disclosed may be readily utilized as a basis for modifying or designing other structures for carrying out the same purposes of the present invention. It should also be realized by those skilled in the art that such equivalent constructions do not depart from the spirit and scope of the invention as set forth in the appended claims.

45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1340
1341
1342
1343
134

BRIEF DESCRIPTION OF THE DRAWING

For a more complete understanding of the present invention, and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawing, in which:

5 FIGURE 1 depicts a multiple source, multiple destination network;

FIGURE 2 depicts an agent device and two separate destination devices wherein dedicated channels link each of the destination devices to the agent device;

10 FIGURE 3 depicts a transaction scheme employing a destination memory control device according to a preferred embodiment of the present invention;

FIGURE 4 depicts a transmission of a first of four transaction messages according to a preferred embodiment of the present invention;

FIGURE 5 depicts a transmission of a second of four transaction messages according to a preferred embodiment of the present invention;

15 FIGURE 6 depicts a transmission of a third of four transaction messages according to a preferred embodiment of the present invention;

FIGURE 7 depicts a transmission of a fourth of four transaction messages according to a preferred embodiment of the present invention; and

20 FIGURE 8 depicts a network having a plurality of source devices and destination devices the communication between which is routed over an unordered, unreliable network according to a preferred embodiment of the present invention.

DETAILED DESCRIPTION

FIGURE 3 depicts a transaction scheme employing a destination memory control device 305 according to a preferred embodiment of the present invention. In a preferred embodiment, source agent 302 is associated with processor 301. Requirements database 303 preferably cooperates with source agent 302 to store the status of transactions and transaction data packets. Network 304 is generally a multiple source, multiple destination computer network but may alternatively be a single device, such as a router, interposed between a single source device and a single destination device. In another alternative embodiment, network 304 may be a local area network or wide area network.

In a preferred embodiment, Content Addressable Memory (CAM) 306 is a memory device which allows any one of its entries to be written to by an external device. CAM 306 preferably includes a feature wherein, when presented with the identification of a source device, or other external device, and an identification of a transaction of interest, memory control 305 will compare the presented transaction identification to all current valid entries in the CAM 306 to determine whether the pertinent transaction is present in the CAM. CAM 306 preferably operates as a database for storing information regarding currently active transactions in network 300.

In a preferred embodiment, processor 301 operates to conduct a transaction with memory board 307, which transaction may be a memory read or memory write operation. The source agent 302 and database 303 at the source site and memory control 305 and CAM 306 at the destination site preferably operate to provide management of the transaction and preservation of data on the progress and status of the transaction so as to minimize latency, avoid duplication of transaction steps, and enable system 300 to possess unambiguous information regarding the status of all transactions pending in system 300 at any given time.

Alternatively, the source 302 and destination 305 devices need not be a processor and memory, respectively, but may include a range of data storage and manipulation devices. For example, a processor to processor transaction, involving a read, write, or other operation is suitable for application of the inventive mechanism. In another alternative embodiment, both

memory boards and a processors may be either sources or destinations in the configuration of FIGURE 3.

In a preferred embodiment, a sequence of four transaction messages or packets are communicated between source agent 302 and memory control 305 in order to establish a status for each transaction at appropriate points within system 300. It will be appreciated that an alternate number of transactions less than or more than four may be employed to perform the functions of the inventive mechanism, and all such variations are within the scope of the present invention. Preferably, the protocol of the present invention ensures that although individual transaction packets may be issued or received more than once, the transaction being implemented by the various transaction packets is performed only once.

FIGURE 4 depicts a transmission of a first of four transaction messages according to a preferred embodiment of the present invention. Reference is made to both FIGURES 3 and 4 in the following discussion. Source agent 302 preferably sends reservation request 307 via network 304 to memory control 305. Reservation request 307 preferably includes information describing a transaction being requested by source agent 302 on behalf of processor chip 301 as well a transaction identification to identify the transaction to memory control 305. Reservation request 307 preferably further includes a time stamp 402 to identify to any device within network 304 encountered by request 307 the time at which reservation request 307 was transmitted by agent 302. The information incorporated in time stamp 402 preferably enables reservation request 402 to be dropped from the system, by any device within the system, upon expiration of a predetermined period (the packet time-out period) beginning with the time indicated in time stamp 402.

In a preferred embodiment, agent 302 includes a timing function 401 for recording the time elapsed since the transmission of reservation request 307. Preferably, agent 302 records, in database 303, the fact and time of transmission of reservation request 307. The transmission of reservation request 307 initiates timer 401 in agent 302 to determine whether reservation response 308 is received within a predetermined period of transmission of reservation request 307. If the source agent time-out period elapses without receipt of

reservation response 308 at agent 302, agent 302 retransmits reservation request 307. The time-out period for dropping reservation request 307 from the system is preferably smaller than the time period established in agent 302 for retransmitting the reservation request, thereby ensuring that any prior reservation request packets have been dropped from the system prior to transmission of a duplicate reservation request. In this manner, the inventive mechanism preferably operates to ensure that there will more than one reservation request pertaining to the same transaction traveling through network 304 at the same time.

In a preferred embodiment, when destination memory control 305 receives reservation request 307, memory control 305 preferably determines whether there is a pre-existing entry in CAM 306 corresponding to reservation request 307. If no corresponding entry is found in CAM 306, memory control 305 preferably concludes that reservation request 307 is being received at memory control 305 for the first time. In this case, the transaction associated with reservation request 307 will preferably be acted upon by writing data to, or reading data from, memory 307.

In a preferred embodiment, if an entry is found in CAM 306 which corresponds to reservation request 307, memory control 306 preferably concludes that reservation request 307 is the result of a resend operation. Generally, at least two options are available to memory control 305 in this case. Where the reservation response generated in response to the packet of which reservation request 307 is a duplicate, is completely stored in CAM 306, control 305 retransmits the stored reservation response as reservation response 308 responsive to reservation request 307. Where, for example, the pertinent transaction is a read operation and data from the read operation is not stored in CAM 306, control 305 will acquire the data from memory board 307 corresponding to the pertinent read operation, and transmit this data as part of reservation response 308 to source 302.

In a preferred embodiment, data resulting from a read operation is not automatically stored in CAM 306. However, where a read transaction is still active in the network, and a memory location which is read from in that read operation will be erased by a prospective write operation or other computing procedure, control 305 preferably loads data from the

memory location (which is soon to be over-written) into the CAM so that it will be available in the event that one or more subsequent reservation requests associated with the pertinent read transaction are received by control 305.

FIGURE 5 depicts a transmission of a second of four transaction messages according to a preferred embodiment of the present invention. Reference is made to both FIGURES 3 and 5 in the following discussion. FIGURE 5 depicts the transmission of reservation response 308 by memory control 305 via network 304 to agent 302. Preferably, generating a reservation response 308 responsive to reservation request 307 includes the steps of generating the data associated with the response, combining this response data with an appropriate transaction identifier to form a reservation response packet, and finally, transmitting reservation response 308.

In a preferred embodiment, the reservation response data is transaction dependent. In the case of a memory “read” operation, the reservation response data preferably comprises the data acquired from memory or memory board 307 as a result of a “read” transaction included in reservation request 307. In the case of a memory “write” operation, there is generally no data to return to source agent 302. The reservation response responsive to a write operation will therefore preferably include an acknowledgment that a requested write operation was successful. Once the reservation response has been prepared, it is stored in the CAM 306 and preferably remains there until a release request packet is received by the memory control 305.

In a preferred embodiment, where a reservation response responsive to the reservation request is not present in the CAM, reservation response 308 is generated and transmitted to agent 302 via network 304 as previously discussed. Where a reservation response corresponding to the reservation request is present in the CAM, indicating that a duplicate reservation request has been received, the already stored reservation response is preferably transmitted to agent 302. Preferably, minimal computational effort is required to transmit this already existing reservation response.

After memory control 305 either discovers reservation response 308 in the CAM 306 or generates the information required therefor, memory control 305 preferably transmits

reservation response to source agent 302 as discussed in connection with FIGURE 5. For the sake of this discussion, it will be assumed that source agent 302 receives reservation response intact and prior to expiration of timer 401 in source agent 302. Timer 401 preferably counts the time elapsed since the transmission of the reservation request 307 to which reservation response 308 is responsive. The system response to reservation request packets and reservation response packets which have been dropped by the system is discussed later in this application.

5 In a preferred embodiment, receipt of reservation response 308 by agent 302 preferably informs agent 302 of the status of the transaction which was included in the reservation request originally transmitted by agent 302. Preferably, reservation response 308 will indicate that the transaction read or write operation was successfully executed. Alternatively however, reservation response 308 may include information indicating that a memory read or write operation was not successful. Further, memory control 305 may optionally insert into reservation response 308, details pertaining to an unsuccessful transaction and available options for correcting the problem. In a preferred embodiment, reservation response 308 returns data to source agent 302 and eventually to processor 302 resulting from execution of a successful transaction.

10 In a preferred embodiment, source agent 302, having received reservation response 308, preferably operates to inform memory control 305 that the data stored in CAM 306 pertaining to the transaction associated with reservation request 307 may be deleted so as to release space in the CAM for use with a new reservation request. It will be appreciated that the released CAM space may be used for a reservation request emerging from any processor connected to network 304 and is not limited to storing such requests only from source agent 302.

20 FIGURE 6 depicts a transmission of a third of four transaction messages according to a preferred embodiment of the present invention. Reference is made to both FIGURES 3 and 6 in the following discussion. FIGURE 6 depicts the transmission of release request 309 from source agent 302 across network 304.

In a preferred embodiment, the transmission of release request by agent 302 signifies that agent 302 and processor 301 regard the transaction transmitted by reservation request 307 as having been concluded, and that stored information pertaining to the completed transaction at a destination site, which in this case includes CAM memory 306, is no longer needed. Accordingly, upon receipt of release request 309 by memory control 305, memory control 305 preferably clears the entry in CAM 306 pertaining to the pertinent transaction.

In a preferred embodiment, as was the case with the transmission of reservation request 307, timer 401 in source agent 302 operates to count elapsed time from the transmission of release request 309. If release response 310 is received before timer 401 reaches a retransmission threshold point, the timer is preferably reset and no retransmission occurs. However, if the timer threshold point is reached prior to receipt of release response 310 at source agent 302, source agent 302 preferably retransmits release request 309. Without any other action, a retransmission operation might leave open the possibility of having two identical release requests in transmission through the system at the same time.

Therefore, in a preferred embodiment, a timing mechanism is implemented in the system to cause a release request to “time out” and drop out of the system after a predetermined period has expired since transmission of that release request. As was described in connection with the reservation request, a time stamp 601 is preferably included in release request 309 which may be read by a number of devices which may be encountered by the release request as it propagates through network 304. Preferably, any one of these devices is enabled to read to keep current time and remove release request 309 from the system if the release request packet time-out period has expired.

In a preferred embodiment, in order to ensure that the release request “times out” before source agent 302 issues a retransmission request, the time-out threshold at which system devices drop the release request from the system is preferably shorter than the time-out threshold which source agent 302 employs to signal retransmission of the release request.

FIGURE 7 depicts a transmission of a fourth of four transaction messages according to a preferred embodiment of the present invention. Reference is made to both FIGURES 3

and 7 in the following discussion. Preferably, memory control 305 transmits release response 310 across network 304 to agent 302.

In a preferred embodiment, transmission of release response 310 by memory control 305 indicates that the transaction which is the subject of the preferably four transaction 5 messages has been completed and that the CAM entry for this transaction has been released, thereby freeing up space in the CAM 306 for a new reservation request. Preferably, upon receipt of release response 310, source agent 302 appropriately concludes that the transaction has been successfully concluded and proceeds to delete an entry corresponding to the completed transaction in database 303.

10 In a preferred embodiment, source agent 302 receives the desired results of a transaction when reservation response 308 is delivered to source agent 302. Source agent 302 then preferably communicates this information in a timely manner to processor 301. Subsequently, the third and fourth transaction messages, the release request and release 15 response, respectively, constitute follow-up operations to appropriately inform memory controller 305 and source agent 302 of the status of the transaction at their respective locations and to free up temporarily allocated storage space in database 303 (on the source agent end) and in CAM 306 (on the destination memory control end). Employing source agent 302 for the management of transaction messages and entering and deleting entries in database 303 preferably beneficially enables processor 301 to concentrate on performing 20 operations on the received data, instead of being burdened with communication and database management tasks.

25 In a preferred embodiment, the provision of services by devices other than the processor, including deleting database and CAM entries for completed transactions and the related communication in both directions of status information, increases processor efficiency and reduces processing and communication latency. The improvement (reduction) in processing latency provided by delegating tasks to the agent 302 and memory control 305 would generally not have been available employing the sequence number systems of the prior art, where processing time would have been spent allocating channels and establishing

sequence numbering systems on both ends of a source-destination pair. The inventive approach therefore provides a significant latency advantage over the systems of the prior art.

Generally, packets transmitted across an unreliable network, such as network 304 may be dropped or corrupted. In a preferred embodiment, the problem of network unreliability and the attendant possibility of packets being dropped is preferably handled by employing time-outs at the source of both reservation requests and release requests as well as inserting a time stamp in both of the reservation request packet 307 and release request packet 309. The interaction time-outs at source agent 302 and time-outs occurring at dispersed locations within the system employing the time stamp in the packets has already been discussed in detail and will therefore not be repeated in detail in this section. While the timer in source agent 302 triggers retransmission of a packet, the combination of the time stamp and the ability of dispersed devices within network 304 to check the current time against a time stamp and time-out period included in a packet, operates to drop out of the system, packets which have "timed out." Preferably, the time-out threshold of the packets traveling through the system is set to be lower than the time out employed by the agent 302 to trigger retransmission of a previously sent packet.

In a preferred embodiment, source agent 302 employs a timer to determine when to retransmit a message where an expected response has not been received. Source agent 302 also includes a time stamp in both the reservation request 307 and release request 309 in order to enable the system to drop these requests from the system if they have not reached the destination memory control 305 within an pre-determined time. In a preferred embodiment, memory control 305 does not perform a similar timing operation for messages transmitted to agent 302. Furthermore, memory control 305 preferably does not include time stamps in reservation response 308 and release response 310. It will be appreciated however, that in an alternative embodiment, the process of counting the time elapsed since a message was transmitted could be performed by memory control 305 alone, or by both source agent 302 and memory control 305. Furthermore, in an alternative embodiment, memory control 305

could include time stamps in the packets for reservation response 308 and release response 310, and all such variations are included within the scope of the present invention.

In a preferred embodiment, agent 302 retransmits packets when its logic indicates that a packet responsive to the transmitted packet should have, but has not been, received after a predefined time period has elapsed since a prior transmission of the packet outbound from the source or agent. Source agent 302 preferably retransmits the relevant packet assuming that the previously transmitted packet has been dropped from the system. Generally, the effect on the system of a dropped packet depends upon which of the four transaction communication packets was dropped.

In a preferred embodiment, if a reservation request is dropped en route to memory control 305, agent 302 will time out according to its schedule, and preferably retransmit a duplicate reservation request by accessing database 303. It will be appreciated that agent 302 is not limited to sending the reservation request only twice, but rather, may repeat the transmission as many times as is necessary to ensure completion of this first transaction packet transmission. However, while transient problems in system 300 may account for the loss of one or two packets, agent 302 may establish a threshold number of failed transmissions above which it regards the network as containing a permanent failure and works to arrange for an alternative method of accomplishing the transmission of the required packet.

In a preferred embodiment, where memory control 305 receives a reservation request representing a transaction not found in the CAM 306, control 305 treats the incoming packet as a first reservation request transmission, preferably establishes a new entry in CAM 306, and preferably produces a reservation response responsive to the reservation request packet. Where memory control 305 receives a reservation request representing a transaction which is found in CAM 306, memory control 305 preferably concludes that a reservation response corresponding to the currently requested transaction was produced and transmitted, but dropped from the system. Accordingly, it is preferably not necessary to generate a new reservation response in response to the duplicate reservation request. Preferably, memory

control 305 generates a duplicate reservation response with information stored by the original reservation response. In a preferred embodiment, the above logic applies in similar manner to the transmission by source agent 302 of release request 309 and the packet responsive thereto, release response 310. In this case, the logic is generally reversed. Since release 5 request 309 generally causes memory control 305 to delete the entry in the CAM 306 for the pertinent transaction, upon receipt of a first release request 309 without any dropped packets, there should not be an entry corresponding to the transaction associated with the release request.

Accordingly, in a preferred embodiment, if a release request 309 is received which does have an entry in the CAM 306, memory control 305 preferably considers the release request 309 to be a first transmission, appropriately acts to delete the pertinent entry in CAM 306, and generates a release response 310. If a release request packet was dropped before reaching memory control 305, memory control 305 preferably interprets the release request packet to be a first transmission, as described above.

Where a release request is received which does not match an entry in CAM 306, memory control preferably considers release request 309 to be a duplicate transmission, and generates and transmits release response 310 to agent 302 across network 304. It will be appreciated that where the release request does not match any entry in CAM 306, there is no entry in CAM 306 which is suitable for deletion. Memory control 305 then preferably generates and transmits a release response 310 to agent 302 across network 304. The sequence of events described above is consistent with a prior release response 310 having been dropped before reaching agent 302.

For purposes of further discussion, the processor and memory are generally referred to herein as "master" devices, and the source agent and destination memory control as "servant" 25 devices. It will be appreciated that the inventive principles are not limited to the master and servant devices being hardware components. The master devices could alternatively be software programs which delegate selected tasks to "servant" programs, or subroutines, so as to remove the delegated tasks from their respective critical paths and thereby execute more

5

efficiently. Likewise, the servant devices may be programs or sub-routines which serve main programs and operate to perform tasks which relieve their respective main programs of certain time consuming tasks. Preferably, the master and servant software components would implement a similar protocol for preserving transaction results and storing transaction packet message statuses as is described above in connection with hardware versions of the master and servant devices.

10
15
20
25

FIGURE 8 depicts a network having a plurality of source and destination devices, the communication between which is routed over an unordered, unreliable network according to a preferred embodiment of the present invention. FIGURE 8 depicts agent 803 associated with processor 801 and agent 804 associated with processor 802. The sequence of dots in between agents 803 and 804 is intended to indicate the presence of a plurality of additional agents also having associated processors. Similarly, memory controls 805 and 806 are associated, respectively, with memory boards 807 and 808. The various agents and memory controls are generally all connected to network 304. In the following paragraph, the combination of an agent and a processor, when acting to request a transaction from a remote location on the network 800, is generally referred to as a source site, and the combination of a memory control and a memory board when operating to service a transaction request received from a remote source site, is generally referred to as a destination site.

20

25

In a preferred embodiment, a transaction requested by a source, executed at a destination, and implemented via a sequence of preferably four transaction packets or messages, as described in connection with FIGURES 3-7 may generally be accomplished between any source-destination combination connected across network 800. For example, where there are one hundred source sites and one hundred destination sites, generally ten thousand different source-destination pairs are available. Although one embodiment of the source-destination pairing may include a source which includes a processor and a destination which includes a memory board, it will be appreciated that source sites are not limited to having processors, and that destination sites are not limited to having memories or memory boards. In alternative embodiments, both processors along with associated equipment, and

memory boards along associated equipment, may operate as either source sites or destination sites, and all such variations are included within the scope of the present invention.

Although the present invention and its advantages have been described in detail, it should be understood that various changes, substitutions and alterations can be made herein without departing from the spirit and scope of the invention as defined by the appended claims. Moreover, the scope of the present application is not intended to be limited to the particular embodiments of the process, machine, manufacture, composition of matter, means, methods and steps described in the specification. As one of ordinary skill in the art will readily appreciate from the disclosure of the present invention, processes, machines, manufacture, compositions of matter, means, methods, or steps, presently existing or later to be developed that perform substantially the same function or achieve substantially the same result as the corresponding embodiments described herein may be utilized according to the present invention. Accordingly, the appended claims are intended to include within their scope such processes, machines, manufacture, compositions of matter, means, methods, or steps.

WHAT IS CLAIMED IS:

1. A method for executing a transaction in a network having a source site and a destination site, the method comprising the steps of:

transmitting an initial transaction request message from said source site to said destination site;

receiving said transaction request message at said destination site;

establishing an association of a data entry with said transaction in a destination database; and

preserving said association of said data entry with said transaction in said destination database so long as said transaction is active in said network.

2. The method of claim 1, comprising the further step of:

executing said transaction at said destination site, thereby producing transaction results.

3. The method of claim 2, comprising the further step of:

where a prospective operation will override said transaction results in a memory board,

storing said transaction results in said destination database, thereby enabling retransmission of said transaction results if a further reservation request message is received at said destination site.

4. The method of claim 2, comprising the further step of:

transmitting said transaction results to said source site over said network.

5. The method of claim 1 comprising the further step of:

transmitting another transaction request message if no response is received from said destination site at said source site within a source site time-out period.

6. The method of claim 5, comprising the further step of:

deleting said initial transaction request message from the network if said transaction request message does not reach said destination site within a request message time-out period, wherein said source site time-out period exceeds said request message time-out period to prevent having two transaction request messages simultaneously in transmission through said network.

7. The method of claim 4, comprising the further steps of:

upon receiving a duplicate transaction request message,
identifying the data entry in the destination database established for said transaction;
acquiring said transaction results; and
retransmitting said acquired transaction results to said source site.

8. The method of claim 7, wherein the step of acquiring comprises:

retrieving said transaction results from said destination database.

9. The method of claim 7, wherein the step of acquiring comprises:

executing said transaction in response to said duplicate transaction request message,
thereby producing said transaction results.

10. The method of step 4, comprising the further steps of:
receiving said transmitted transaction results at said source site; and
transmitting, from said source site to said destination site, a release request to delete
said data entry associated with said transaction in said destination database.

11. The method of step 10, comprising the further steps of:
receiving at said destination site, said release request to delete said data entry
associated with said transaction; and
deleting, within said destination database, said data entry associated with said
transaction, thereby liberating space in said destination database.

12. The method of step 11, comprising the further step of:
transmitting, from said destination site to said source site, a release response message,
thereby indicating that said data entry associated with said transaction in said destination
database has been deleted.

13. The method of step 1, comprising the further step of:
wherein the source site includes a processor and an agent device,
delegating said step of transmitting said initial transaction request message to said
agent device.

14 A system for reliably executing a transaction at a destination site requested by a source site, the system comprising:

means for transmitting an initial transaction request message to said destination site from said source site;

5 means for executing a transaction associated with said initial transaction request message at said destination site;

a reservation database at said destination site for storing information uniquely identifying said transaction.

15. The system of claim 14, wherein the reservation database is a content addressable memory.

16. The system of claim 14, wherein the source site comprises:

a processor; and the destination site comprises:

a memory.

17. The system of claim 16, wherein the source site further comprises:

a processor agent device for conducting communication with said destination site, thereby enabling said processor to efficiently concentrate on other tasks.

18. The system of claim 17, wherein the source site further comprises:

a source site database for preserving an identification and a status of said transaction until said transaction is complete.

19. The system of claim 16, wherein the processor agent device comprises:
a timer for initiating a retransmission of said transaction request message if no
message responsive to said initial transaction request message is received at said processor
agent device upon expiration of a retransmission time-out period.

20 A system for executing a transaction in a network having a source site and a destination site, the system comprising:

means for transmitting an initial transaction request message from said source site to said destination site;

5 means for receiving said transaction request message at said destination site;

means for establishing an association of a data entry with said transaction in a destination database; and

means for preserving said association of said data entry with said transaction in said destination database so long as said transaction is active in said network.

**PROTOCOL FOR INSURING EXACTLY ONCE SEMANTICS OF
TRANSACTIONS ACROSS AND UNORDERED, UNRELIABLE NETWORK**

ABSTRACT OF THE DISCLOSURE

The present invention implements a mechanism for reliably communicating transaction messages between source and destination devices connected across an unreliable network, where a transaction is an operation requested by a source device of a destination device and where a plurality of messages are exchanged between source and destination in order to ensure completion of the transaction. Because the network linking the source and destination is unreliable, the source and destination operate to preserve data regarding the transaction messages which they have received and transmitted. If responsive messages are not timely received, they are generally resent. A dual timing system is preferably implemented to ensure that there are never two identical messages simultaneously in transmission through the network.

5

10

3/1/00

809273.1




FIG. 1




FIG. 2




FIG. 3








FIG. 8



**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**

ATTORNEY DOCKET NO. 10992150-1

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

PROTOCOL FOR INSURING EXACTLY ONCE SEMANTICS OF TRANSACTIONS ACROSS AN UNORDERED, UNRELIABLE NETWORK

the specification of which is attached hereto unless the following box is checked:

(was filed on _____ as US Application Serial No. or PCT International Application Number _____ and was amended on _____ (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119	
N/A			YES: _____	NO: _____
			YES: _____	NO: _____

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION SERIAL NUMBER	FILING DATE
N/A	

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION SERIAL NUMBER	FILING DATE	STATUS (patented/pending/abandoned)
N/A		

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879

Place Customer
Number Bar Code
Label here

Send Correspondence to:
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80528-9599

Direct Telephone Calls To:

Alexander J. Neudeck
(970) 898-4931

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Tony M. Brewer Citizenship: U.S.A.

Residence: 5227 Mariners Drive, Plano, Texas 75093

Post Office Address: 5227 Mariners Drive, Plano, Texas 75093

Tony M. Brewer
Inventor's Signature

Date

2128/00

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)**

ATTORNEY DOCKET NO. 10992150-1

Full Name of # 2 joint inventor: Gregory S. Palmer Citizenship: U.S.A.

Residence: 3012 Mason Drive, Plano, Texas 75025

Post Office Address: 3012 Mason Drive, Plano, Texas 75025

Gregory S. Palmer Date 2/28/00
Inventor's Signature Date

Full Name of # 3 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 4 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 5 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 6 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 7 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 8 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____